

CLAIMS

What is claimed is:

1. A method of operating a storage system, comprising:
storing a plurality of data blocks having a horizontal redundant relationship;
5 storing a plurality of checksums, each checksum having a vertical redundant relationship
with a corresponding one of the plurality of data blocks; and
in response to detection of a data integrity error in at least one of the plurality of data
blocks, reviewing the vertical redundant relationships between each of the
checksums in the plurality of checksums and the corresponding data blocks.
- 10 2. The method of Claim 1, wherein
said reviewing the vertical redundant relationships comprises calculating a new checksum
for each of the data blocks in the plurality of data blocks and comparing the
calculated new checksum to the checksum from the plurality of checksums
15 corresponding to that data block.
3. The method of Claim 1, wherein:
said storing the plurality of checksums comprises storing the plurality of checksums such
that each data block in the plurality of data blocks has a vertical redundant
20 relationship with a set of corresponding checksums, and each of the checksums in
the set of corresponding checksums has a horizontal checksum redundant
relationship with the other checksums in the set of corresponding checksums.
4. The method of Claim 3, further comprising:

in response to detecting a first data integrity error between a first data block and a first checksum from a first set of checksums corresponding to the first data block, reviewing the vertical redundant relationship between the first data block and the remaining checksums in the first set of checksums; and

5 in response to detecting no data integrity errors between the first data block and the remaining checksums in the first set of checksums, creating a restored first checksum to replace the first checksum corresponding to the first data integrity error.

10 5. The method of Claim 1, further comprising:

in response to detection of a data integrity error in at least one of the plurality of data blocks, reviewing the horizontal redundant relationship between the data block having the data integrity error and the remaining data blocks in the plurality of data blocks.

15 6. The method of Claim 5, wherein:

said plurality of data blocks comprises a plurality of data stripe units and a parity stripe unit; and

said reviewing the horizontal redundant relationship comprises calculating a new parity stripe unit based on the plurality of data stripe units and comparing the calculated new parity stripe unit to the parity stripe unit from the plurality of data blocks.

20 7. The method of Claim 5, wherein:

said plurality of data blocks comprises a first mirrored data block and a corresponding second mirrored data block; and

Sub
e2

said reviewing the horizontal redundant relationship comprises comparing the first mirrored data block and the second mirrored data block.

8. The method of Claim 5, further comprising:

5 in response to detecting a first data integrity error in exactly one of the vertical redundant relationships and detecting no data integrity errors in the horizontal redundant relationship, creating a restored checksum using the data block corresponding to the first data integrity error to replace the existing checksum corresponding to the first data integrity error.

10 9. The method of Claim 5, further comprising:

15 in response to detecting a first data integrity error in exactly one of the vertical redundant relationships and detecting a second data integrity error in the horizontal redundant relationship, creating a restored data block using the horizontal redundant relationship to replace the data block corresponding to the first data integrity error and comparing the restored data block to the checksum corresponding to the first data integrity error.

20 10. The method of Claim 9, further comprising:

in response to a mismatch between the restored data block and the checksum corresponding to the first data integrity error, creating a restored checksum using the restored data block to replace the checksum corresponding to the first data integrity error.

25 11. The method of Claim 5, further comprising:

Sub
ar

in response to detecting a first data integrity error in a first vertical redundant relationship, detecting a second data integrity error in a second vertical redundant relationship, and detecting no data integrity errors in the horizontal redundant relationship, reporting a data loss condition.

5

12. The method of Claim 5, further comprising:

in response to detecting a first data integrity error in a first vertical redundant relationship, detecting a second data integrity error in a second vertical redundant relationship, and detecting a third data integrity error in the horizontal redundant relationship:

10

creating a first restored data block using the horizontal redundant relationship to replace the data block corresponding to the first vertical redundant relationship;

comparing the first restored data block to the checksum corresponding to the first vertical redundant relationship;

15

creating a second restored data block using the horizontal redundant relationship to replace the data block corresponding to the second vertical redundant relationship; and

comparing the second restored data block to the checksum corresponding to the second vertical redundant relationship.

20

13. The method of Claim 12, further comprising:

reporting a data loss condition in response to either:

(a) a mismatch between the first restored data block and the checksum corresponding to the first vertical redundant relationship and a mismatch

25

Sub
a2

between the second restored data block and the checksum corresponding to the second vertical redundant relationship; or

- 5 (b) a match between the first restored data block and the checksum corresponding to the first vertical redundant relationship and a match between the second restored data block and the checksum corresponding to the second vertical redundant relationship.

14. The method of Claim 12, further comprising:

10 in response to a match between the first restored data block and the checksum corresponding to the first vertical redundant relationship and a mismatch between the second restored data block and the checksum corresponding to the second vertical redundant relationship, creating a restored checksum using the second restored data block to replace the checksum corresponding to the second vertical redundant relationship.

- 15 15. The method of Claim 1, further comprising:

in response to detecting data integrity errors in at least three of the vertical redundant relationships, reporting a data loss.

- 20 16. A storage system, comprising:

a storage array comprising a plurality of storage devices;

a plurality of data blocks having a horizontal redundant relationship, each data block stored on one of the plurality of storage devices;

25 a plurality of checksums, each checksum having a vertical redundant relationship with a corresponding one of the plurality of data blocks; and

an array manager coupled to the storage array, wherein said array manager is configured such that in response to detection of a data integrity error in at least one of the plurality of data blocks, said array manager reviews the vertical redundant relationships between each of the checksums in the plurality of checksums and the corresponding data blocks.

17. The storage system of Claim 16, wherein:

said array manager is configured to review the vertical redundant relationships by calculating a new checksum for each of the data blocks in the plurality of data blocks and comparing the calculated new checksum to the checksum from the plurality of checksums corresponding to that data block.

18. The storage system of Claim 16, wherein:

each data block in the plurality of data blocks has a vertical redundant relationship with a set of corresponding checksums, and each of the checksums in the set of corresponding checksums has a horizontal checksum redundant relationship with the other checksums in the set of corresponding checksums.

19. The storage system of Claim 18, wherein:

said array manager is configured such that in response to detecting a first data integrity error between a first data block and a first checksum from a first set of checksums corresponding to the first data block, said array manager reviews the vertical redundant relationship between the first data block and the remaining checksums in the first set of checksums; and

said array manager is further configured such that in response to detecting no data integrity errors between the first data block and the remaining checksums in the

first set of checksums, said array manager creates a restored first checksum to replace the first checksum corresponding to the first data integrity error.

20. The storage system of Claim 16, wherein:

said array manager is configured such that in response to detection of a data integrity error in at least one of the plurality of data blocks, said array manager reviews the horizontal redundant relationship between the data block having the data integrity error and the remaining data blocks in the plurality of data blocks.

21. The storage system of Claim 20, wherein:

said plurality of data blocks comprises a plurality of data stripe units and a parity stripe unit; and

said array manager is configured to review the horizontal redundant relationship by calculating a new parity stripe unit based on the plurality of data stripe units and comparing the calculated new parity stripe unit to the parity stripe unit from the plurality of data blocks.

22. The storage system of Claim 20, wherein:

said plurality of data blocks comprises a first mirrored data block and a corresponding second mirrored data block; and

said array manager is configured to review the horizontal redundant relationship by comparing the first mirrored data block and the second mirrored data block.

23. The storage system of Claim 20, wherein:

5 said array manager is configured such that in response to detecting a first data integrity error in exactly one of the vertical redundant relationships and detecting no data integrity errors in the horizontal redundant relationship, said array manager creates a restored checksum using the data block corresponding to the first data integrity error to replace the existing checksum corresponding to the first data integrity error.

24. The storage system of Claim 20, wherein:

10 said array manager is configured such that in response to detecting a first data integrity error in exactly one of the vertical redundant relationships and detecting a second data integrity error in the horizontal redundant relationship, said array manager creates a restored data block using the horizontal redundant relationship to replace the data block corresponding to the first data integrity error and compares the restored data block to the checksum corresponding to the first data integrity error.

15 25. The storage system of Claim 24, wherein:

20 said array manager is configured such that in response to a mismatch between the restored data block and the checksum corresponding to the first data integrity error, said array manager creates a restored checksum using the restored data block to replace the checksum corresponding to the first data integrity error.

26. The storage system of Claim 20, further comprising:

25 said array manager is configured such that in response to detecting a first data integrity error in a first vertical redundant relationship, detecting a second data integrity error in a second vertical redundant relationship, and detecting no data integrity errors in the horizontal redundant relationship, said array manager reports a data loss condition.

Sub
ad

27. The storage system of Claim 20, wherein:

5 said array manager is configured such that in response to detecting a first data integrity error in a first vertical redundant relationship, detecting a second data integrity error in a second vertical redundant relationship, and detecting a third data integrity error in the horizontal redundant relationship, said array manager:

creates a first restored data block using the horizontal redundant relationship to replace the data block corresponding to the first vertical redundant relationship;

10 compares the first restored data block to the checksum corresponding to the first vertical redundant relationship;

creates a second restored data block using the horizontal redundant relationship to replace the data block corresponding to the second vertical redundant relationship; and

15 compares the second restored data block to the checksum corresponding to the second vertical redundant relationship.

28. The storage system of Claim 27, wherein:

20 said array manager is configured such that said array manager reports a data loss condition in response to either:

(a) a mismatch between the first restored data block and the checksum corresponding to the first vertical redundant relationship and a mismatch between the second restored data block and the checksum corresponding to the second vertical redundant relationship; or

25 (b) a match between the first restored data block and the checksum corresponding to the first vertical redundant relationship and a match between the second

restored data block and the checksum corresponding to the second vertical redundant relationship.

29. The storage system of Claim 27, wherein:

said array manager is configured such that in response to a match between the first restored data block and the checksum corresponding to the first vertical redundant relationship and a mismatch between the second restored data block and the checksum corresponding to the second vertical redundant relationship, said array manager creates a restored checksum using the second restored data block to replace the checksum corresponding to the second vertical redundant relationship.

30. The storage system of Claim 16, wherein:

said array manager is configured to report a data loss condition in response to detecting data integrity errors in at least three of the vertical redundant relationships.

31. A computer-readable medium whose contents cause a computer system to operate a storage system by performing the steps of:

detecting a data integrity error in at least one of a plurality of data block having a horizontal redundant relationship; and

analyzing a plurality of checksums, each checksum having a vertical redundant relationship with a corresponding one of the plurality of data blocks, said analyzing comprising reviewing the vertical redundant relationships between each of the checksums in the plurality of checksums and the corresponding data blocks.

32. The computer-readable medium of Claim 31, wherein:

Sub
a2

said reviewing the vertical redundant relationships comprises calculating a new checksum for each of the data blocks in the plurality of data blocks and comparing the calculated new checksum to the checksum from the plurality of checksums corresponding to that data block.

5

33. The computer-readable medium of Claim 31, wherein:

each data block in the plurality of data blocks has a vertical redundant relationship with a set of corresponding checksums, and each of the checksums in the set of corresponding checksums has a horizontal checksum redundant relationship with the other checksums in the set of corresponding checksums.

10

34. The computer-readable medium of Claim 33, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

15

in response to detecting a first data integrity error between a first data block and a first checksum from a first set of checksums corresponding to the first data block, reviewing the vertical redundant relationship between the first data block and the remaining checksums in the first set of checksums; and

20

in response to detecting no data integrity errors between the first data block and the remaining checksums in the first set of checksums, creating a restored first checksum to replace the first checksum corresponding to the first data integrity error.

25

35. The computer-readable medium of Claim 31, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

in response to detection of a data integrity error in at least one of the plurality of data blocks, reviewing the horizontal redundant relationship between the data block having the data integrity error and the remaining data blocks in the plurality of data blocks.

5

36. The computer-readable medium of Claim 35, wherein:

said plurality of data blocks comprises a plurality of data stripe units and a parity stripe unit; and

said reviewing the horizontal redundant relationship comprises calculating a new parity stripe unit based on the plurality of data stripe units and comparing the calculated new parity stripe unit to the parity stripe unit from the plurality of data blocks.

10

37. The computer-readable medium of Claim 35, wherein:

said plurality of data blocks comprises a first mirrored data block and a corresponding second mirrored data block; and

said reviewing the horizontal redundant relationship comprises comparing the first mirrored data block and the second mirrored data block.

15

38. The computer-readable medium of Claim 35, wherein the contents of the computer-

readable medium cause the computer system to operate the storage system by performing the further steps of:

20

in response to detecting a first data integrity error in exactly one of the vertical redundant relationships and detecting no data integrity errors in the horizontal redundant relationship, creating a restored checksum using the data block corresponding to the first data integrity error to replace the existing checksum corresponding to the first data integrity error.

25

39. The computer-readable medium of Claim 35, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

in response to detecting a first data integrity error in exactly one of the vertical redundant relationships and detecting a second data integrity error in the horizontal redundant relationship, creating a restored data block using the horizontal redundant relationship to replace the data block corresponding to the first data integrity error and comparing the restored data block to the checksum corresponding to the first data integrity error.

40. The computer-readable medium of Claim 39, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

in response to a mismatch between the restored data block and the checksum corresponding to the first data integrity error, creating a restored checksum using the restored data block to replace the checksum corresponding to the first data integrity error.

41. The computer-readable medium of Claim 35, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

in response to detecting a first data integrity error in a first vertical redundant relationship, detecting a second data integrity error in a second vertical redundant relationship, and detecting no data integrity errors in the horizontal redundant relationship, reporting a data loss condition.

42. The computer-readable medium of Claim 35, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

in response to detecting a first data integrity error in a first vertical redundant

relationship, detecting a second data integrity error in a second vertical redundant relationship, and detecting a third data integrity error in the horizontal redundant relationship:

creating a first restored data block using the horizontal redundant relationship to replace the data block corresponding to the first vertical redundant relationship;

comparing the first restored data block to the checksum corresponding to the first vertical redundant relationship;

creating a second restored data block using the horizontal redundant relationship to replace the data block corresponding to the second vertical redundant relationship; and

comparing the second restored data block to the checksum corresponding to the second vertical redundant relationship.

43. The computer-readable medium of Claim 42, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

reporting a data loss condition in response to either:

- (a) a mismatch between the first restored data block and the checksum corresponding to the first vertical redundant relationship and a mismatch between the second restored data block and the checksum corresponding to the second vertical redundant relationship; or

Sub
a2

(b) a match between the first restored data block and the checksum corresponding to the first vertical redundant relationship and a match between the second restored data block and the checksum corresponding to the second vertical redundant relationship.

5

44. The computer-readable medium of Claim 42, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

in response to a match between the first restored data block and the checksum

10 corresponding to the first vertical redundant relationship and a mismatch between the second restored data block and the checksum corresponding to the second vertical redundant relationship, creating a restored checksum using the second restored data block to replace the checksum corresponding to the second vertical redundant relationship.

15

45. The computer-readable medium of Claim 31, wherein the contents of the computer-readable medium cause the computer system to operate the storage system by performing the further steps of:

in response to detecting data integrity errors in at least three of the vertical redundant relationships, reporting a data loss.

20

46. A storage system, comprising:

a plurality of storage devices configured to store a plurality of data blocks and a plurality of checksums, wherein each data block in said plurality of data blocks has a first type of redundant relationship with a corresponding one of the plurality of checksums and each data block in said plurality of data blocks has a second type

25

Sub
a2

of redundant relationship with the remaining data blocks in said plurality of data blocks; and

an array manager coupled to said plurality of storage devices, wherein said array manager is configured to detect a data integrity error in one of the first type of redundant relationships between the plurality of data blocks and the plurality of checksums and to check for additional data integrity errors in the remaining first type of redundant relationships between the plurality of data blocks and the plurality of checksums.

10 47. The storage system of Claim 46, wherein:

15 said array manager is configured to check for additional data integrity errors in the remaining first type of redundant relationships between the plurality of data blocks and the plurality of checksums by calculating a new checksum for each of the data blocks in the plurality of data blocks and comparing the calculated new checksum to the checksum from the plurality of checksums corresponding to that data block.

20 48. The storage system of Claim 46, wherein:

each data block in the plurality of data blocks has a vertical redundant relationship with a set of corresponding checksums, and each of the checksums in the set of corresponding checksums has a horizontal checksum redundant relationship with the other checksums in the set of corresponding checksums.

25 49. The storage system of Claim 48, wherein:

said array manager is configured such that in response to detecting a first data integrity error between a first data block and a first checksum from a first set of checksums corresponding to the first data block, said array manager reviews the redundant

relationships between the first data block and the remaining checksums in the first set of checksums; and

said array manager is further configured such that in response to detecting no data integrity errors between the first data block and the remaining checksums in the first set of checksums, said array manager creates a restored first checksum to replace the first checksum corresponding to the first data integrity error.

50. The storage system of Claim 46, wherein:

said array manager is configured such that in response to detection of a data integrity error in at least one of the plurality of data blocks, said array manager reviews the second type of redundant relationship between the data block having the data integrity error and the remaining data blocks in the plurality of data blocks.

51. The storage system of Claim 50, wherein:

said plurality of data blocks comprises a plurality of data stripe units and a parity stripe unit; and

said array manager is configured to review the second type of redundant relationship by calculating a new parity stripe unit based on the plurality of data stripe units and comparing the calculated new parity stripe unit to the parity stripe unit from the plurality of data blocks.

52. The storage system of Claim 50, wherein:

said plurality of data blocks comprises a first mirrored data block and a corresponding second mirrored data block; and

said array manager is configured to review the second type of redundant relationship by comparing the first mirrored data block and the second mirrored data block.

53. The storage system of Claim 50, wherein:

said array manager is configured such that in response to detecting a first data integrity error in exactly one of the first type of redundant relationships and detecting no data integrity errors in the second type of redundant relationship, said array manager creates a restored checksum using the data block corresponding to the first data integrity error to replace the existing checksum corresponding to the first data integrity error.

54. The storage system of Claim 50, wherein:

said array manager is configured such that in response to detecting a first data integrity error in exactly one of the first type of redundant relationships and detecting a second data integrity error in the second type of redundant relationship, said array manager creates a restored data block using the second type of redundant relationship to replace the data block corresponding to the first data integrity error and compares the restored data block to the checksum corresponding to the first data integrity error.

55. The storage system of Claim 54, wherein:

said array manager is configured such that in response to a mismatch between the restored data block and the checksum corresponding to the first data integrity error, said array manager creates a restored checksum using the restored data block to replace the checksum corresponding to the first data integrity error.

56. The storage system of Claim 50, further comprising:

5 said array manager is configured such that in response to detecting a first data integrity error in a first one of the first type of redundant relationship, detecting a second data integrity error in a second one of the first type of redundant relationship, and detecting no data integrity errors in the second type of redundant relationship, said array manager reports a data loss condition.

57. The storage system of Claim 50, wherein:

10 said array manager is configured such that in response to detecting a first data integrity error in a first one of the first type of redundant relationship, detecting a second data integrity error in a second one of the first type of redundant relationship, and detecting a third data integrity error in the second type of redundant relationship, said array manager:

15 creates a first restored data block using the second type of redundant relationship to replace the data block corresponding to the first one of the first type of redundant relationship;

compares the first restored data block to the checksum corresponding to the first one of the first type of redundant relationship;

20 creates a second restored data block using the second type of redundant relationship to replace the data block corresponding to the second one of the first type of redundant relationship; and

compares the second restored data block to the checksum corresponding to the second one of the first type of redundant relationship.

58. The storage system of Claim 57, wherein:

25 said array manager is configured such that said array manager reports a data loss condition in response to either:

- 5
- (a) a mismatch between the first restored data block and the checksum corresponding to the first one of the first type of redundant relationship and a mismatch between the second restored data block and the checksum corresponding to the second one of the first type of redundant relationship; or
- (b) a match between the first restored data block and the checksum corresponding to the first one of the first type of redundant relationship and a match between the second restored data block and the checksum corresponding to the second one of the first type of redundant relationship.
- 10

59. The storage system of Claim 57, wherein:

15

said array manager is configured such that in response to a match between the first restored data block and the checksum corresponding to the first one of the first type of redundant relationship and a mismatch between the second restored data block and the checksum corresponding to the second one of the first type of redundant relationship, said array manager creates a restored checksum using the second restored data block to replace the checksum corresponding to the second one of the first type of redundant relationship.

20 60. The storage system of Claim 46, wherein:

said array manager is configured to report a data loss condition in response to detecting data integrity errors in at least three of the first type of redundant relationships.